

1 SUMMARY OF THE INVENTION

2 Thus, in one aspect of the present invention, there is now provided methods, systems and
3 apparatus for transferring interrupts from a peripheral device to a host computer system.
4 In an example embodiment, the apparatus comprising: a buffer for storing indications of
5 interrupts generated by the peripheral device; and a controller for, in response to a preset
6 condition being met, generating a control data block having a payload portion, moving
7 the contents of the buffer to the payload portion of the control data block, and sending the
8 control data block to the host computer system. The buffer preferably comprises a first in
9 - first out memory buffer

10 In some embodiments, the present invention also extends to a peripheral device
11 comprising apparatus as herein before described and to a data communications network
12 interface comprising such a peripheral device. The present invention further extends to a
13 data processing system comprising a host processing system having a memory, a data
14 communications interface for communicating data between the host computer system and
15 a data communications network, and apparatus as hereinbefore described for controlling
16 flow of interrupts from the data communication interface to the memory of the host
17 computer system.

18 In another aspect of the present invention , there is now provided a method for
19 transferring interrupts from a peripheral device to a host computer system, the method
20 comprising: storing interrupts generated by the peripheral device in a buffer; determining
21 if a preset condition is met, and, in response to the preset condition being met, generating
22 a control data block having a payload portion, moving the contents of the buffer to the
23 payload portion of the control data block, and sending the control data block to the host
24 computer system.

1 BRIEF DESCRIPTION OF THE DRAWINGS

2 These and other objects, features, and advantages of the present invention will become
3 apparent upon further consideration of the following detailed description of the invention,
4 by way of example only, when read in conjunction with the drawing figures, in which:

5 Figure 1 is a block diagram of an example of a data processing network;

6 Figure 2 is a block diagram of a network interface adapter card for the data processing
7 network;

8 Figure 3 is a block diagram of an example of a host computer system for the data
9 network;

10 Figure 4 is a block diagram of an example of an Integrated System on a Chip (ISOC) for
11 the network adapter card;

12 Figure 5 is another block diagram of the ISOC;

13 Figure 6 is a block diagram of the ISOC demonstrating information flow through the
14 ISOC ;

15 Figure 7 is a block diagram of a logical transmit path through the ISOC;

16 Figure 8 is a block diagram of a logical receive path through the ISOC;

17 Figure 9A is a block diagram of a cyclic descriptor table

18 Figure 9B is a block diagram of a linked set of descriptor tables;

- 1 Figure 10 is a block diagram of a virtual buffer and its physical counterpart buffer;
- 2 Figure 11 is a block diagram of a completion queue;
- 3 Figure 12 is a block diagram of a transmit flow of data from the host to the network;
- 4 Figure 13 is another block diagram of a transmit flow of data from the host to the
5 network;
- 6 Figure 14 is a block diagram of a receive flow of data from the network to the host;
- 7 Figure 15 is another block diagram of a receive flow of data from the network to the host.
- 8 Figure 16 is another block diagram of the ISOC;
- 9 Figure 17 is of an interrupt flow between the ISOC and the host computer system; and
- 10 Figure 18 is a block diagram of an Interrupt Control Block.

11 DESCRIPTION OF THE INVENTION

12 The present invention provides methods, systems and apparatus for transferring interrupts
13 from a peripheral device to a host computer system. An example of apparatus comprises:
14 a buffer for storing indications of interrupts generated by the peripheral device; and a
15 controller for, in response to a preset condition being met, generating a control data block
16 having a payload portion, moving the contents of the buffer to the payload portion of the

1 control data block, and sending the control data block to the host computer system. The
2 buffer preferably comprises a first in - first out memory buffer

3 Preferably, the preset condition comprises a determination that the buffer is full. The
4 preset condition may comprise a determination that at least a predetermined plurality of
5 indications is stored in the buffer and that a predetermined period has elapsed. Similarly,
6 the preset condition may comprise a determination that at least one indication is stored in
7 the buffer and that a predetermined period has elapsed.

8 In some embodiments of the present invention, the control data block comprises a header
9 portion having an identifier for identifying the ICB and a count indicative of the number
10 of indications included in the payload portion. The header portion may also comprise a
11 time of day stamp.

12 The present invention extends to a peripheral device comprising apparatus as herein
13 before described, and to a data communications network interface, comprising such a
14 peripheral device. The present invention also extends to a data processing system
15 comprising a host processing system having a memory, a data communications interface
16 for communicating data between the host computer system and a data communications
17 network, and apparatus as hereinbefore described for controlling flow of interrupts from
18 the data communication interface to the memory of the host computer system.

19 The present invention further provides a method for transferring interrupts from a
20 peripheral device to a host computer system. In an example embodiment, the method
21 comprising: storing interrupts generated by the peripheral device in a buffer; determining
22 if a preset condition is met, and, in response to the preset condition being met, generating
23 a control data block having a payload portion, moving the contents of the buffer to the
24 payload portion of the control data block, and sending the control data block to the host
25 computer system.

1 Referring first to Figure 1, an example of a data processing network embodying the
2 present invention includes a plurality of host computer systems 10, and a plurality of
3 attached devices 20, interconnected by an intervening network architecture 30 such as an
4 InfiniBand network architecture (InfiniBand is a trade mark of the InfiniBand Trade
5 Association). The network architecture 30 typically comprises a plurality of data
6 communications switches 40. The host computer systems 10 and the attached devices 20
7 each form a node in the data processing network. Each host computer system 10 includes
8 a plurality of central processing units (CPUs) 50, and a memory 60 interconnected by a
9 bus architecture 70 such as a PCI bus architecture. A network adapter 80 is also
10 connected to the bus architecture for communicating data between the host computer
11 system 10 and other nodes in the data processing network via the network architecture 30.

12 Referring now to Figure 2, in more particular embodiments of the present invention, the
13 network adapter 80 comprises a pluggable option card having a connector, such as an
14 edge connector, for removable insertion into the bus architecture 70 of the host computer
15 system 10. The option card carries an Application Specific Integrated Circuit (ASIC) or
16 Integrated System on a Chip (ISOC) 120 connectable to the bus architecture 70 via the
17 connector 170, one or more third level memory modules 250 connected to the ISOC 120,
18 and an interposer 260 connected to the ISOC 120 for communicating data between the
19 media of the network architecture 30 and the ISOC 120. The interposer 260 provides a
20 physical connection to the network 30.

21 In some embodiments of the present invention, the interposer 260 may be implemented in
22 a single ASIC. However, in other embodiments of the present invention, the interposer
23 260 may be implemented by multiple components. For example, if the network 30
24 comprises an optical network, the interposer 260 may comprise a retimer driving a
25 separate optical transceiver. The memory 250 may be implemented by SRAM, SDRAM,
26 or a combination thereof. Other forms of memory may also be employed in the

1 implementation of memory 250. The ISOC 120 includes a first and a second memory.
2 The memory subsystem of the adapter 80 will be described shortly.

3 As will become apparent from the following description, this arrangement provides:
4 improved performance of distributed applications operating on the data processing
5 network; improved system scalability; compatibility with a range of communication
6 protocols; and reduced processing requirements in the host computer system. More
7 specifically, this arrangement permits coexistence of heterogeneous communication
8 protocols between the adapters 80 and the host systems 10. Such protocols can serve
9 various applications, use the same adapter 80, and use a predefined set of data structures
10 thereby enhancing data transfers between the host and the adapter 80. The number of
11 application channels that can be opened in parallel is determined by the amount of
12 memory resources allocated to the adapter 80 and is independent of processing power
13 embedded in the adapter. It will be appreciated from the following that the ISOC 120
14 concept of integrating multiple components into a single integrated circuit chip
15 component advantageously minimizes manufacturing costs and in provides reusable
16 system building blocks. However, it will also be appreciated that in other embodiments
17 of the present invention, the elements of the ISOC 120 may be implemented by discrete
18 components.

19 In the following description, the term Frame refers to data units or messages transferred
20 between software running on the host computer system 10 and the adapter 80. Each
21 Frame comprises a Frame Header and a data payload. The data payload may contain user
22 data, high level protocol header data, acknowledgments, flow control or any combination
23 thereof. The contents of the Frame Header will be described in detail shortly. The
24 adapter 80 processes only the Frame Header. The adapter 80 may fragment Frames into
25 smaller packets which are more efficiently transported on the network architecture 30.
26 However, such fragmentation generally does not transform the data payload.

1 In other particular embodiments of the present invention, data is transported on the
2 network architecture 30 in atomic units hereinafter referred to as Packets. Each Packet
3 comprises route information followed by hardware header data and payload data. In a
4 typical example of the present invention, a packet size of up to 1024 bytes is employed.
5 Frames of larger size are fragmented into 1024 byte packets. It will be appreciated that in
6 other embodiments of the present invention, different packet sizes may be employed.

7 In some embodiments of the present invention, communications between the adapter 80
8 and multiple applications running on the host computer system 10 are effected via a
9 Logical Communication Port architecture (LCP). The adapter 80 comprises a memory
10 hierarchy which allows optimization of access latency to different internal data structures.
11 This memory hierarchy will be described shortly. In preferred embodiments of the
12 present invention, the adapter 80 provides separate paths for outbound (TX) data destined
13 for the network architecture 30 and inbound (RX) data destined for the host computer
14 system 10. Each path includes its own data transfer engine, header processing logic and
15 network architecture interface. These paths will also be described in detail shortly.

16 Referring now to Figure 3, the LCP architecture defines a framework for the interface
17 between local consumers running on the host computer system 10 and the adapter 80.
18 Examples of such consumers include both applications and threads. The computer
19 system 10 can be subdivided into a user application space 90 and a kernel space 110. The
20 LCP architecture provides each consumer with a logical port into the network architecture
21 30. This port can be accessed directly from a user space 90. In particularly preferred
22 embodiments of the present invention, a hardware protection mechanism takes care of
23 access permission. An LCP registration is performed by the kernel space 110 prior to
24 transfer of data frames. The LCP architecture need not define a communication protocol.
25 Rather, it defines an interface between the applications and the adapter 80 for transfer of
26 data and control information. Communication protocol details may be instead set by the
27 application and program code executing in the adapter 80. The number of channels that

1 can be used on the adapter 80 is limited only by the amount of memory on the adapter
2 card 80 available for LCP related information. Each LCP port can be programmable to
3 have a specific set of features. The set of features is selected according to the specific
4 protocol to best support data transfer between the memory 60 in the host computer system
5 and the adapter 80. Various communication protocols can be supported simultaneously,
6 with each protocol using a different LCP port.

7 The LCP architecture comprises LCP Clients 100, an LCP Manager 130 resident in the
8 kernel space 130, and one or more LCP Contexts 140 resident in the adapter 80. Each
9 LCP Client 100 is a unidirectional application end point connected to an LCP port. An
10 LCP client 100 can be located in the user application space 90 or in the kernel 110. In
11 operation, each LCP client 100 produces commands and data to be read from the memory
12 60 and transferred by the adapter 80 via a TX LCP channel, or consumes data transferred
13 by the adapter 80 to the memory 60 via an RX LCP channel. The LCP Manager 130 is a
14 trusted component that services request for LCP channel allocations and deallocations
15 and for registration of read/write areas in the memory 60 for each channel. The LCP
16 Manager 130 allows a user space application to use resources of the adapter 80 without
17 compromising other communication operations, applications, or the operating system of
18 the host computer system 10.

19 Each LCP Context 140 is the set of control information required by the adapter 80 to
20 service a specific LCP Client 100. The LCP Context 140 may include LCP channel
21 attributes which are constant throughout existence of the channel, such as possible
22 commands, pointer structure, and buffer descriptor definitions. The LCP Context 140
23 may also include specific LCP service information for the LCP channel, such as the
24 amount of data waiting for service, and the next address to access for the related LCP
25 channel. The LCP context 140 is stored in memory resident in the adapter 80 to enable
26 fast LCP context switching when the adapter 80 stops servicing one channel and starts
27 servicing another channel.

1 An LCP Client 100 requiring initiation of an LCP port turns to the LCP Manager 130 and
2 requests the allocation of an LCP channel. The LCP channel attributes are determined at
3 this time and prescribe the behavior of the LCP port and the operations that the LCP
4 Client 100 is authorized to perform in association with the LCP port. The LCP Client
5 100 is granted an address that will be used to access the adapter 80 in a unique and secure
6 way. This address is known as a Doorbell Address.

7 The LCP Manager 130 is also responsible for registering areas of the host memory 60 to
8 enable virtual to physical address translation by the adapter, and to allow user space
9 clients to access these host memory areas without tampering with other programs.
10 Registration of new buffers and deregistration of previous buffers can be requested by
11 each LCP Client 100 during run-time. Such a change, requires a sequence of information
12 exchanges between the LCP Client 100, the LCP Manager 130, and the adapter 80.
13 Each LCP Client 100 and port are associated with an LCP Context 140 that provides all
14 the information required by the adapter 80 to service pending requests sent by the LCP
15 port for command execution.

16 To initiate memory transfers between the LCP Client 100 and the adapter 80, and initiate
17 transmission of frames, the LCP Client 100 prepares descriptors holding the information
18 for a specific operation. The LCP Client 100 then performs an I/O write to the Doorbell
19 address mapped to the adapter 80. Writing to the Doorbell address updates the LCP
20 Context 140 on the adapter 80, adding the new request for execution. The adapter 80
21 arbitrates between various transmit LCP ports that have pending requests, and selects the
22 next one to be serviced. On receipt of data, the Frame and LCP for a received packet are
23 identified. Descriptors are generated to define the operation required for the receive LCP.
24 Execution of these descriptors by an LCP Engine of the adapter 80 stores the incoming
25 data in an appropriate data buffer allocated to the LCP channel in the memory 60 of the
26 host computer system 10. For each LCP channel serviced, the adapter 80 loads the

1 associated LCP context information and uses this information to perform the desired set
2 of data transfers. The adapter 80 then continues on to process the next selected LCP
3 Context 140.

4 Referring now to Figure 3, and as mentioned earlier, the ISOC 120 comprises a first
5 memory space 220 and 230 and a second memory space 240 and the adapter 80 further
6 comprises a third level memory 250. The first, second, and third memory spaces for part
7 of a memory subsystem 210 of the adapter 80. In a preferred embodiment of the present
8 invention, the ISOC 120 comprises a TX processor (TX MPC) 150 dedicated to data
9 transmission operations and an RX processor (RX MPC) 160 dedicated to data reception
10 operation. In particularly preferred embodiments of the present invention, processors 150
11 and 160 are implemented by Reduced Instruction Set Computing (RISC) microprocessors
12 such as IBM PowerPC 405 RISC microprocessors. Within the memory subsystem 210,
13 the ISOC 120 comprises, in addition to the first and second memory spaces, a data cache
14 180 and an instruction cache 170 associated with TX processor 150, together with a
15 second data cache 190 and second instruction cache 190 associated with RX processor
16 160. The difference between the three levels is the size of memory and the associated
17 access time. As will become apparent shortly, the memory subsystem 210 facilitates:
18 convenient access to instruction and data by both the TX processor 150 and the RX
19 processor 160; scalability; and sharing of resources between the TX processor 150 and
20 the RX processor 160 in the interests of reducing manufacturing costs.

21 The first level memory spaces (M1) 220 and 230 comprise a TX-M1 memory space 220
22 and RX-M1 memory space 230. The TX-M1 memory 220 can be accessed only by the
23 TX processor 150 and the RX-M1 memory 230 can be accessed only by the RX processor
24 160. In operation the first level memory spaces 220 and 230 are used to hold temporary
25 data structures, header templates, stacks, etc. The first level memory spaces 220 and 230
26 both react to zero wait states. Each one of the first level memory spaces 220 and 230 is
27 connected only to the data interface of the corresponding one of the processors 150 and

1 160 and not to the instruction interface. This arrangement enables both cacheable and
2 non-cacheable first level memory areas available while maintaining efficient access to
3 data in the first level memory spaces 230 and 240.

4 The second level memory space (M2) 240 is a shared memory available to both
5 processors 150 and 160, other components of the adapter 80, and to the host computer
6 system 10. Access to the second level memory space 240 is slower than access to the first
7 level memory areas 220 and 230 because the second level memory space 240 is used by
8 more agent via a shared internal bus. The third level memory space 250 is also a shared
9 resource. In particularly preferred embodiments of the present invention the adapter 80
10 comprises a computer peripheral circuit card on which the first level memory spaces 220
11 and 230 and the second level memory space 240 are both integrated on the same ASIC as
12 the processors 150 and 160. The shared memory spaces 240 and 250 are generally used
13 for data types that do not require fast and frequent access cycles. Such data types include
14 LCP contexts 140 and virtual address translation tables. The shared memory spaces 240
15 and 250 are accessible to both instruction and data interfaces of the processors 150 and
16 160.

17 The adapter 80 handles transmission and reception data flows separately. The separate
18 processor 150 and 160 for the transmission and reception path avoids the overhead of
19 switching between task, isolates temporary processing loads in one path from the other
20 path, and facilitates use of two embedded processors to process incoming and outgoing
21 data streams. Referring now to Figure 5, the ISOC 120 comprises transmission path logic
22 280 and reception path logic 290, and shared logic 300. The transmission path logic 280
23 comprises an LCP TX engine 310 for decoding specifics of each LCP channel and
24 fetching LCP related commands for execution; TX logic 320 for controlling transfer of
25 frames into the adapter 80, the aforementioned TX processor 150 for managing TX frame
26 and packet processing; the aforementioned first level TX memory 220 for holding
27 instructions and temporary data structures; and link logic 330; and logic for assisting the

1 TX processor 150 in managing the data flow and packet processing such as routing
2 processing for fragmentation of frames into data packets. The TX processor 150
3 processes tasks in series based on a polling only scheme in which the processor is
4 interrupted only on exceptions and errors. The first level TX memory 220 is employed by
5 the processor 150 for communicating with TX logic 320. The reception path logic 290
6 comprises link logic 340; hardware for assisting the aforementioned RX processor 160 in
7 processing headers of incoming packets and transformation or assembly of such packets
8 into frames; the aforementioned RX processor 160 for RX frame and packet processing;
9 the aforementioned first level RX memory 230 for holding instructions; RX logic 350 for
10 controlling transfer of frames from the network architecture 30; and an LCP RX engine
11 360 for decoding the specifics of each LCP channel, storing the incoming data in the
12 related LCP data structures in the memory 60 of the host computer system, and accepting
13 and registering pointers to empty frame buffers as they are provided by the LCP Client
14 100 for use by the adapter 80. The RX processor 160 processes tasks in series using a
15 polling only scheme in which the RX processor 160 is interrupted only on exceptions or
16 errors. The level 1 RX memory 230 is used by the RX processor 160 to communicate
17 with the RX logic 350.

18 As mentioned earlier, the ISOC approach permits reduction in manufacturing costs
19 associated with the adapter 80 and the other components thereof, such as the circuit board
20 and the other supporting modules. The ISOC approach also increases simplicity of the
21 adapter 80, thereby increasing reliability. The number of connections between elements
22 of the ISOC 120 is effectively unlimited. Therefore, multiple and wide interconnect paths
23 can be implemented. In the interests of reducing data processing overheads in the host
24 computer system 10, data transfer operations to and from the host memory 60 are
25 predominantly performed by the ISOC 120. The ISOC 120 also performs processing of
26 the header of incoming and outgoing packets. During transmission, the ISOC 120 builds
27 the header and routes it to the network architecture 30. During reception, the adapter 80
28 processes the header in order to determine its location in the system's memory. The level

1 1 memories 220 and 230 are zero wait state memories providing processor data space
2 such as stack, templates, tables, and temporary storage locations. In especially preferred
3 embodiments of the present invention, the transmission path logic 280, reception path
4 logic 290, and shared logic 300 are built from smaller logic elements referred to as cores.
5 The term core is used because there elements are designed as individual pieces of logic
6 which have stand-alone properties enabling them to be used for different applications.

7 As indicated earlier, the transmission path logic 280 is responsible for processing
8 transmission or outgoing frames. Frame transmission is initiated via the bus architecture
9 70 by a CPU such as CPU 50 of the host computer system 10. The ISOC 120 comprises
10 bus interface logic 370 for communicating with the bus architecture 70. The ISOC 120
11 also comprises bus bridging logic 390 connecting the bus interface logic 370 to a
12 processor local bus (PLB) 390 of the ISOC 120. The TX LCP engine 310 fetches
13 commands and frames from the host memory 60. The TX processor 150 processes the
14 header of each frame into a format suitable for transmission as packets on the network
15 architecture 30. The TX logic 320 transfer the frame data without modification. The link
16 logic 330 processes each packet to be transmitted into a final form for transmission on the
17 network architecture 30. The link logic 330 may comprises one or more ports each
18 connectable to the network architecture 30.

19 As indicated earlier, the reception path logic 290 is responsible for processing incoming
20 packets. Initially, packets received from the network architecture 30 are processed by
21 link logic 340. Link logic 340 recreates the packet in a header and payload format. To
22 determine the packet format and its destination in the host memory 60, the header is
23 processing by the RX processor 230. The link logic 340 may comprises one or more
24 ports each connectable to the network architecture 30. The RX LCP engine is responsible
25 for transferring the data into the host memory 60 via the bus architecture 70.

1 The transmission path logic 280 comprises a HeaderIn first in- first out memory (FIFO)
2 400 between the TX LCP engine 310 and the TX processor 220. The reception path logic
3 comprises a HeaderOut FIFO 410 between the RX processor 230 and the RX LCP engine
4 360. Additional FIFOs and queues are provided in the TX logic 320 and the RX logic
5 350. These FIFOs and queues will be described shortly.

6 The shared logic 300 comprises all logical elements shared by the transmission path logic
7 280 and the reception path logic 290. These elements include the aforementioned bus
8 interface logic 370, bus bridging logic 380, PLB 390, second level memory 240 and a
9 controller 420 for providing access to the remote third level memory 250. The bus
10 interface logic 370 operates as both master and slave on the bus architecture 70. As a
11 slave, the bus interface logic allows the CPU 50 to access the second level memory 240,
12 the third level memory 250 via the controller 420, and also configuration registers and
13 status registers of the ISOC 120. Such registers can generally be accessed by the CPU 50,
14 the TX processor 150 and the RX processor 160. As a master, the bus interface logic
15 allows the TX LCP engine 310 and the RX LCP engine 360 to access the memory 60 of
16 the host computer system 10. In Figure 5, “M” denotes a master connection and “S”
17 denotes a slave connection.

18 Referring now to Figure 6, packet flow through the ISOC 120 is generally symmetrical.
19 In other words, the general structure of flow is similar in both transmit and receive
20 directions. The ISOC 120 can be regarded as comprising first interface logic 440; a first
21 control logic 460; processor logic 480; second control logic 470; and second interface
22 logic 450. Packets are processed in the following manner:

23 A. In the transmit direction, information is brought into the ISOC 120 from the bus
24 architecture 70 through the first interface logic. In the receive direction,
25 information is brought into the ISOC 120 from the network architecture 30
26 through the second interface logic 450.

- 1 B. In the transmit direction, information brought into the ISOC 120 through the first
2 interface logic 440 is processed by the first control logic 460. In the receive
3 direction, information brought into the ISOC through the second interface logic
4 450 is processed by the second control logic 470.
- 5 C. In the transmit direction, a frame header is extracted for an outgoing frame at the
6 first control logic 460 and processed by the processor logic 480. The processor
7 logic 480 generates instructions for the second control logic 470 based on the
8 frame header. The payload of the outgoing frame is passed to the second interface
9 logic 470. In the receive direction, a frame header is extracted from an incoming
10 frame at the second control logic 470 and processed by the processor logic 480.
11 The processor logic 480 generates instructions for the first control logic 460 based
12 on the frame header. The payload of the incoming frame is passed to the first
13 control logic 460. In both directions, the processor 480 is not directly handling
14 payload data.
- 15 D. In the transmit direction, the second control logic 470 packages the outgoing
16 payload data according to the instructions received from the processor logic 480.
17 In the receive direction, the first control logic 460 packages the incoming payload
18 according to the instructions received from the processor logic 480.
- 19 E. In the transmit direction, the information is moved through the second interface
20 logic 450 to its destination via the network architecture 30. In the receive
21 direction, the information is moved through the first interface logic to its
22 destination via the bus architecture 70.

1 An interface to software operating on the host computer system 10 is shown at 430.
2 Similarly, interfaces to microcode operating on the processor inputs and outputs is shown
3 at 490 and 500.

4 Referring to Figure 7, what follows now is a more detailed description of one example of
5 a flow of transmit data frames through the ISOC 120. The ISOC 120 can be divided into
6 an LCP context domain 510, a frame domain 520 and a network domain 530 based on the
7 various formats of information within the ISOC 120. The TX LCP engine 310 comprises
8 an LCP requests FIFO 550, Direct Memory Access (DMA) logic 560, frame logic 580,
9 and the aforementioned LCP context logic 140. The LCP request FIFO 550, DMA logic
10 560, and LCP TX Context logic 590 reside in the LCP context domain 510. The frame
11 logic 580 resides in the frame domain 520. The TX logic 320, first level TX memory
12 space 220, and TX processor 150 straddle the boundary between the frame domain 520
13 and the network domain 530. The TX link logic 330 resides in the network domain 530.
14 In particularly preferred embodiments of the present invention, the HeaderIn FIFO 400 is
15 integral to the first level TX memory space 220. In general, an application executing on
16 the host computer system 10 creates a frame. The frame is then transmitted using a TX
17 LCP channel on the adapter 80. Handshaking between the application and the adapter 80
18 assumes a prior initialization performed by the LCP Manager 130. To add an LCP
19 Service Request, an LCP Client 100 informs the adapter 80 that one or more additional
20 transmit frames are ready to be executed. This is performed by writing to a control word
21 in to a Doorbell. The Doorbell's addresses are allocated in such as way that the write
22 operation is translated into a physical write cycle on the bus architecture 70, using an
23 address that is uniquely associated with the LCP port and protected from access by other
24 processes. The adapter 80 detects the write operation and logs the new request by
25 incrementing an entry of previous requests for the specific LCP Client 100. This is part
26 of the related LCP Context 140. An arbitration list, retained in the memory subsystem
27 210 of the adapter 80 is also updated. In a simple example, arbitration uses the
28 aforementioned FIFO scheme 550 between all transmit LCP channels having pending

1 requests. While one LCP channel is serviced, the next LCP channel is selected. The
2 service cycle begins when the corresponding LCP Context is loaded into the TX LCP
3 engine 310. The LCP Context 140 is then accessed to derive atomic operations for
4 servicing the LCP channel and to determine parameters for such operations. For
5 example, such atomic operations may be based on LCP channel attributes recorded in the
6 LCP Context 140. A complete service cycle typically includes a set of activities
7 performed by the adapter 80 to fetch and execute a plurality of atomic descriptors created
8 by the LCP Client 100. In the case of a TX LCP channel, the service cycle generally
9 includes reading multiple frames from the host memory 60 into the memory subsystem
10 210 of the adapter 80. Upon conclusion, all the LCP Context information requiring
11 modification (in other words, the LCP Service Information) is updated in the memory
12 subsystem 210 of the adapter 80. In general, the first action performed by the adapter 80
13 within the LCP Service cycle, is to fetch the next descriptor to be processed.

14 Processing of transmission frames by the ISOC 120 typically includes the following
15 steps:

16 A. Fetching the subsequent LCP port frame descriptor.

17 The address of the next descriptor to be fetched is stored as parts of the LCP
18 channel's Context 140. The adapter 80 reads the descriptor from host memory 60
19 and decodes the descriptor based on the LCP channel attributes. The descriptor
20 defines the size of the new frame header, the size of the data payload, and the
21 location of these items.

22 B. Conversion of virtual address to physical address.

23 If a data buffer is referenced by virtual memory addresses in an application, the
24 address should go through an additional process of address translation. In this

1 case, the virtual address used by the application is translated into a physical
2 address usable by the adapter 80 while it access the host memory 60. This is done
3 by monitoring page boundary crossings and using physical page location
4 information written by the LCP manager 130 into the memory subsystem 210 of
5 the adapter 80. The virtual to physical translation process serves also as a security
6 measure in cases where a descriptor table is created by an LCP client 100 which is
7 not trusted. This prevents unauthorized access to unrelated areas of the host
8 memory 60.

9 C. Reading the frame header.

10 Using physical addressing, the header and payload data of the TX frame are read
11 from buffers in the host memory 60. The header is then stored in the TX
12 HeaderIn FIFO 400. When the header fetch is completed, the adapter 80 sets an
13 internal flag indicating that processing of the header can be initiated by the TX
14 processor 150.

15 D. Reading the frame data.

16 The payload data is read from the host memory 60 and stored by the adapter 80 in
17 a data FIFO 570. The data FIFO 570 is shown in Figure 7 as resident in the TX
18 logic 320. However, the data FIFO 570 may also be integral to the first level TX
19 memory space 220. Data read transactions continue until all data to be
20 transmitted is stored in the memory subsystem 210 of the adapter 80. Following
21 completion of the read operation, a status indication is returned to the LCP Client
22 100. Note that processing of the header can start as soon as the header has been
23 read into the HeaderIn FIFO 400. There is no need to wait for the whole data to
24 be read.

1 E. Processing the frame header

2 The header processing is performed by the TX processor 150. Header processing
3 is protocol dependent and involves protocol information external to the LCP
4 architecture. The TX processor 150 runs TX protocol header microcode and
5 accesses routing tables and other relevant information already stored in the
6 memory subsystem 210 of the adapter 80 during a protocol and routing
7 initialization sequence. When the TX processor 150 receives an indication that a
8 new header is waiting in the HeaderIn FIFO 400, it starts the header processing.
9 The header processing produces one or more packet headers which are in the
10 format employed to send packets over the network architecture 30 and include
11 routing information. If the payload size is larger than a maximum packet size
12 allowed by the network architecture 30, the payload is fragmented by generating
13 several packet headers each used in connection with consecutive data segments of
14 the original payload data to form packets for communication over the network
15 architecture 30.

16 F. Queuing the packet header for transmission

17 A command defining the number of header words and the number of data words
18 for a packet and the packet header itself are written by the TX processor 150 to a
19 TX HeaderOut FIFO 540 in the first level memory space 220.

20 G. Merging packet header and packet data for transmission.

21 Transmission of a packet on the network architecture 30 is triggered whenever a
22 command is ready in the HeaderOut FIFO 540, and the data FIFO 570 contains
23 enough data to complete the transmission of the related packet. A Cyclic
24 Redundancy Check (CRC) may be added to the header and data of each packet.

1 Each complete packet is transferred to the network architecture 30 via the TX link
2 logic 330.

3 The transmission process for each frame is completed when all the frame data is
4 transmitted on the network architecture 30, by means of one or more packets. For each
5 frame processed by the adapter 80, a status may be returned to the application via a
6 second LCP Client 100. This status indicates the completion of the frame data transfer
7 from the host memory 60 onto the adapter 80, completion of the frame transmission itself,
8 or other levels of transmission status. At any instance in time, the adapter 80 may be
9 concurrently executing some or all of the following actions: selecting the next LCP to be
10 serviced; initiating service for LCP channel A; executing DMA fetch of data for the last
11 frame of LCP channel B; processing a frame header and fragmentation for LCP channel C
12 ; and, transmitting packets originated by LCP channel D.

13 Referring to Figure 8, what follows now, by way of example only, is a description of a
14 data frame reception by an application using an RX LCP port. The operation of the ISOC
15 120 may vary depending on the type of protocol supported by the LCP. Handshaking
16 between the application and the adapter 80 assumes a prior initialization performed by the
17 LCP manager 130. The RX LCP engine 360 comprises LCP allocation logic 620, LCP
18 Context logic 610, and DMA logic 630 all residing in the LCP domain 520. The RX
19 processor 160, first level RX memory space 230, and RX logic 350 all straddle the
20 boundary between the frame domain 520 and the network domain 530. The RX link
21 logic 340 and packet assist logic 600 reside in the network domain 530. In particularly
22 preferred embodiments of the present invention, the HeaderOut FIFO 410 is located in
23 the first level RX memory space 230. Frames received by the ISOC 120 from the
24 network architecture 30 are written into LCP client buffers in the host memory 60.
25 Availability of memory buffers is determined by the LCP RX client 100 and is indicated
26 to the adapter 80 for insertion of incoming data frames. The LCP client 100 provides
27 buffers by writing into a receive Doorbell on the ISOC 120, similar to the aforementioned

1 manner in which the transmission path logic 280 is informed of new frames ready to be
2 transmitted. The Doorbell register address is allocated such that the write operation is
3 translated into a physical write cycle on the bus architecture 70. The adapter 80 detects
4 the write operation and logs the new provision of empty memory areas by incrementing
5 the number of available word entries for the specific LCP RX Client 100. The available
6 word count is part of the related LCP context 140. Whenever an application completes
7 processing of a received frame within a buffer, it writes to the Doorbell. The write cycle
8 indicates the number of words in the newly available memory space. The count within
9 the LCP context is incremented by that amount. A packet received from the network
10 architecture 30 may be part of a larger frame that will be assembled by the adapter 80 into
11 contiguous space in the host memory 60. Processing of received frames by the ISOC 120
12 generally includes the following steps:

13 A. Splitting packet header and data

14 The RX link logic 340 translates information from the network architecture 30
15 into a stream of packets. Each received packet is processed by the RX link logic
16 340 to separate the packet header from the payload data. The header is pushed
17 into an RX HeaderIn FIFO 640 in the first level RX memory space 230. The
18 payload is pushed into an RX data FIFO 650 in the RX logic 350. The RX data
19 FIFO 650 may also be implemented in the first level RX memory space 230.

20 B. Decoding the packet header and generating and LCP frame header.

21 The packet header is decoded to provide fields indicative of an ID for the frame to
22 which the packet belongs, the size of the payload, and the size of the frame data.
23 Once the packet header is reader for the RX HeaderIn FIFO 640, an indication is
24 sent to the RX processor 160. The RX processor processes the packet header
25 information and generates an LCP related command including information

1 required to transfer the packet data. Such information includes packet address and
2 length. At the end of the header processing, a descriptor, or a set of descriptors,
3 are written to the LCP RX HeaderOut FIFO 410, and an indication is triggered.

4 C. Transfer of data within the RX LCP Context.

5 The descriptors are fetched from the RX HeaderOut FIFO 410 by the RX LCP
6 engine 360, and then decoded. The descriptors include the LCP number, packet
7 address, packet data length and the source address of the data to be transferred in
8 the memory subsystem 210 of the adapter 80. The RX LCP engine 340 uses the
9 LCP Context information to create a target physical address (or addresses if a
10 page is crossed) to be written to in the host memory 60 and initiates DMA
11 transfers to write the data.

12 D. ISOC DMA transactions.

13 The ISOC 120 aims to optimize transactions on the bus architecture 70 by
14 selecting appropriate bus commands and performing longest possible bursts.

15 At any instance in time, the adapter 80 may be concurrently executing some or all of the
16 following: processing a buffer allocation for LCP channel X; initiating an inbound data
17 write service for LCP channel A; executing a DMA store of data for LCP channel B;
18 processing a frame assembly of a packet destined for LCP channel C; and, receiving
19 packets for LCP channel D.

20 To minimize frame processing overhead on the RX processor 160 and TX processor 150,
21 packet assist logic 600 sometimes comprises frame fragmentation logic, CRC and
22 checksum calculation logic, and multicast processing logic. The data flow between both
23 the TX and RX LCP engines 310 and 360 and the host 10 will now be described in detail.

1 Both TX and RX LCP ports use memory buffers for transferring data and descriptor
2 structures that point to such memory buffers. The descriptor structures are used to
3 administer data buffers between a data provider and a data consumer and to return empty
4 memory buffers to be used by the data provider. The descriptors point to the memory
5 buffers based on either physical or virtual addresses. TX LCP channels are responsible
6 for data transfer from the host memory 60 into buffers of the ISOC 120. Other layers of
7 logic are responsible for transferring data from buffers of the ISOC 120 into the network
8 30. RX LCP channels are responsible for transferring data received from the network 30
9 to the host memory 60.

10 The TX and RX LCP engines 310 and 360 are capable of handling a relatively large
11 number of LCP channels. Each LCP channel has a set of parameters containing all
12 information specific thereto. The information comprises the configuration of the channel,
13 current state and status. The LCP context 140 associated with a channel is set by the LCP
14 manager 130 during initialization of the channel. During channel operation, the content
15 of the LCP context 140 is updated only by the ISOC 120. The LCP contexts 140 are
16 saved in a context table within the memory subsystem 210 of the adapter 80. Access to
17 the LCP context 140 of an LCP channel is performed according to the LCP number. The
18 LCP RX and TX channels use different LCP context structures.

19 Data buffers are pinned areas in the memory 60 of the host 10. Transmit buffers hold
20 data that for transmission. The TX LCP engine 310 moves the data located in these
21 buffers into internal buffers of the ISOC 120. Incoming data received from the network
22 30 is moved by the RX LCP engine 360 into buffers in the memory 60 of the host 10.
23 Ownership of the buffers alternates between software in the host 10 and the ISOC 120.
24 The order of events on LCP TX channels is as follows:

25 A. Software in the host 10 prepares buffers with data to be transmitted in the memory
26 60 of the host 10;

- 1 B. The software notifies the ISOC 120 that data in the buffers is ready to be
- 2 transmitted;
- 3 C. The ISOC 120 reads the data from the buffers; and,
- 4 D. The ISOC 120 identifies to the software in the host 10 the buffers that were read
- 5 and can be reused by the software in the host 10 to transfer new data.

6 The order of events on LCP RX channels is preferably as follows:

- 7 A. The software in the host 10 prepares buffers into which the ISOC 210 can write
- 8 the received data;
- 9 B. The software notifies the ISOC 120 that free buffers are ready in the memory 60
- 10 of the host;
- 11 C. The ISOC 120 writes the data to the buffers; and,
- 12 D. The ISOC 120 identifies to the software in the host 10 the buffers that were filled
- 13 with received data and can be processed by the software.

14 When the software prepares buffers to be used by the ISOC 120, buffer information is

15 tracked via doorbell registers. Information relating to buffers used by the ISOC 120 is

16 returned to the software using a status update or through a completion queue. For TX

17 LCP channels, the buffers include data and header information transferred by the TX LCP

18 engine 310 into the ISOC 120 and processed to become one or more packets for

19 transmission on the network 30. The header is used by the TX processor 150 of the ISOC

20 120 to generate the header of the packet to be transmitted on the network 30. For RX

21 LCP channels, free buffers are assigned by the software in the host 10 to the adapter 80.

22 The adapter 80 fills the buffers with the received packets.

23 The descriptors have defined data structures known to both the ISOC 120 and software in

24 the host 10. The software uses descriptors to transfer control information to the ISOC

25 120. The control information may be in the form of a frame descriptor, a pointer

1 descriptor, or a branch descriptor depending on desired function. Descriptor logic in the
2 software and in the ISOC 120 generate and modify the descriptors according to control
3 measures to be taken. Such measure will be described shortly. A frame descriptor
4 comprises a description of the packet (e.g.: data length, header length, etc.). A pointer
5 descriptor comprises a description of a data location. A branch descriptor comprises
6 description of the descriptor location (e.g.: link lists of descriptors). Information in the
7 descriptors is used for control by the software in the host 10 of the data movement
8 operations performed by the TX and RX LCP engines 310 and 360. The information
9 used to process a frame to generate a TX packet header is located in the header of the
10 frame. Referring to Figure 9A, descriptors may be provided in a single table 700 with the
11 LCP context 140 pointing to the head of the table 700. Referring to Figure 9B,
12 descriptors may also be arranged in a structure of linked descriptor tables 720-740.
13 Following LCP channel initialization, the LCP context 140 points to the head of the first
14 descriptor table 720 in the structure. Branch descriptors 750-770 are used to generate a
15 linked list of tables 720-740 where a branch descriptor 750-770 at the end of a descriptor
16 table 720-740 points to the beginning of another table 720-740. Referring back to
17 Figure 9A, branch descriptors can also be used to generate a cyclic buffer where a branch
18 descriptor 710 at the end of a table 700 points to the beginning of the same table 700. A
19 cyclic buffer may also be used in the receive path. In this case, the LCP 140 context is
20 initiated to point to the head of the buffer. The buffer is wrapped around when the ISOC
21 120 reaches its end. The software in the host 10 can write the descriptors into the
22 memory 60 in the host 10 (for both the receive and the transmit paths) or into the memory
23 250 of the adapter 80 (for the transmit path only). Writing descriptors to the memory
24 subsystem 210 of the adapter 80 involves an I/O operation by the software in the host 10
25 and occupies the memory subsystem 210 of the adapter 80. Writing descriptors in the
26 memory 60 of the host 80 requires the adapter 80 to access the memory 60 of the host 10
27 whenever it has to read a new descriptor. The location of the software descriptors is
28 defined by the LCP manager 130 for each LCP channel independently. The location of

1 the descriptors is defined according to system performance optimization. The descriptors
2 provide flexibility in the construction of queues.

3 The RX and TX LCP engines 310 and 360 use addresses to access the descriptors in the
4 descriptor tables and to access data buffers. An address can be either a physical address
5 or a virtual address. The term physical address describes an address that the ISOC 120
6 can drive, as is, to the bus 70. The term virtual address describes an address which is not
7 a physical one and is used by the software or microcode. The virtual address has to pass
8 through a mapping in order to generate the physical address. An address used by the TX
9 and RX LCP engines 310 and 360 can have different sources as follows: pointer in the
10 LCP channel context 140; pointer in descriptors prepared by software running on the host
11 10; pointer in descriptors prepared by the RX processor 160; and, pointer in descriptors
12 prepared by the TX processor 150 (used for returning a completion message). A pointer
13 can point to a descriptor or to a data buffer. Every address used by the TX and RX LCP
14 engines 310 and 360 can be optionally mapped to a new address used as the physical
15 address on the bus 70. The address mapping is done by the TX and RX LCP engines 310
16 and 360. The ISOC 120 uses local memory 210 to hold the translation tables. The LCP
17 manager 130 writes the translation tables to the adapter 80 during memory registration.
18 The address mapping allows virtual addressing to be used for buffers or descriptor tables.
19 The virtual addressing enables the management of virtual buffers that are physically
20 located in more than one physical page. The address mapping also allows the host 10 to
21 work directly with applications using virtual addresses without requiring a translation
22 processor for the software.

23 Referring to Figure 10, shown therein is an image 800 of a buffer 880 as it appears to the
24 software in the host 10. Also shown is a physical mapping 810 of the address at it is used
25 to access the memory 60 in the host 10. A virtual pointer points 820 to a location in the
26 buffer. The buffer in this example is a virtual buffer occupying a few noncontiguous
27 pages 840-870 in the memory 60 of the host 10. The LCP engines 310 and 360 perform

1 the mapping by translating the address via a translation table 830. The translation table
2 holds a physical address pointer to the head of each physical buffer 840-870 mapped from
3 the virtual buffer 880. Address mapping in the adapter 80 allows flexibility when
4 mapping descriptors and data buffers in the memory 60 in the host 10. Address mapping
5 in the adapter 80 also allows a direct connection to software buffers that use virtual
6 addresses without requiring the software in the host 10 to perform address translation to a
7 physical address.

8 Each packet which the adapter 80 writes to the memory 60 in the host has a status
9 associated therewith. The status allows synchronization between the adapter 80 and the
10 software in the host 10. The status can be used to indicate different reliability levels of
11 packets. The ISOC 120 provides the following status write backs: Transmit DMA
12 Completion indicates that a data in a TX packet has been read into the adapter 80;
13 Reliable Transmission is returned to indicate the completion of data transmission in the
14 network 30; Receive DMA Completion indicates completion of a receive data transfer
15 into the memory 60; and, Reliable Reception indicates reception of a transmit packet by a
16 destination node in the network 30.

17 A TX frame descriptor includes a 2 byte status field. Status write back means that a
18 transaction status is written back into a descriptor. The status includes a completion bit
19 which can be polled by the software in the host 10. When the software in the host 10
20 finds a set completion bit, it may reuse the buffers associated with the frame defined by
21 the frame descriptor.

22 A completion queue is implemented by an RX LCP channel. The LCP channel used by
23 the completion queue has all the flexibility and properties that can be implemented by any
24 RX LCP channel. The TX and RX processor 150 and 160 generates status write backs to
25 indicate reliable transmission, reliable reception, receive DMA completion, or transmit
26 DMA completion. Different indications relating to the frame are used in different cases.

1 For example, in the case of a reliable transmission, the TX processor 150. Reads internal
2 registers indicating the status of a packet transmission. In the case of reliable reception,
3 the RX processor 160 gets a completion indication as a received packet which includes an
4 acknowledgment. In the case of a receive DMA completion, the RX processor 160 uses
5 frame completion information. In the case of a transmit DMA completion, the TX
6 processor 150 indicates the reception of a frame for transmission in the adapter 80. A
7 completion queue can be used by a single TX or RX LCP channel or may shared by
8 multiple channels. Micro code in the adapter 80 updates a status queue by initiating a
9 frame descriptor into a command queue of the RX LCP engine 360. Referring to Figure
10 11, the status is transferred to the memory 60 of the host 10 via a completion status LCP
11 900 comprising a completion queue 920. The completion queue 900 is continuous (either
12 physically or virtually) and is located in the memory 60 of the host 10. For example, the
13 completion queue can be held in a continuous buffer. Entries 930 in the completion
14 queue preferably have a fixed size. Each entry holds a pointer 940 to the head of a buffer
15 950 associated with a receive LCP 910. The buffer 950 is filled by the packet 960
16 associated with the completion status..

17 A TX software/adaptor handshake comprises an TX LCP port and an completion RX LCP
18 port. Each LCP transmit channel uses the following data structures:

19 A Doorbell entry, implemented as a memory mapped address, informs the adapter
20 80 of incremental requests to process descriptors and data. Each process has a
21 unique access into a single page of memory mapped address used for Doorbell
22 access.

23 An LCP context entry in the adapter memory space 210, containing LCP attributes
24 and status fields.

1 A structure of transmit descriptors. This structure may span across multiple
2 physical pages in the memory 60 of the host 10. If virtual addressing is used for
3 the descriptors, a translation table is used to move one page to the next. If
4 physical addressing is used for the descriptors, branch descriptors are used to
5 move from one page to the next. Transmit descriptors contain a status field that
6 can be updated following transfer of all descriptor related data to the adapter 80.

7 Transmit data buffers pinned in the memory 60 of the host 10 pointed to by the
8 pointer descriptors. If virtual addressing is used for the data buffers, a translation
9 tale converts the pointer into physical addresses used by the adapter 80 to access
10 the memory 60 in the host 10.

11 A translation table and protection blocks in the adapter memory space 210 are
12 used for address mapping.

13 Referring to Figure 12, a transmit packet flow comprises, at step 1000, software 1020 in
14 the host 10 filling buffer 1030 with data to be transmitted. At step 1010, the software
15 1020 updates the descriptors 1040. The descriptors 1040 may be either in the memory 60
16 of the host 10 or in the memory subsystem 210 of the adapter 80. At step 1050, the
17 software 1020 rings the Doorbell to notify the adapter 80 that new data is ready to be
18 transmitted. At step 1060, the adapter 80 manages arbitration between requests from the
19 different LCP channels. When a channel wins the arbitration, the adapter 80 reads the
20 new descriptors 1040. At step 1070, the adapter 80 reads the data. At step 1080, the data
21 is transmitted to the network 30. At step 1090, the status is updated in the descriptors
22 1040 or in the completion queue.

23 The TX LCP channel may use address translation when accessing data buffers. In this
24 case, the data buffer is composed of multiple memory pages. As far as the process is
25 concerned, these memory pages are in consecutive virtual memory space. However, as

1 far as the adapter 80 is concerned, these memory pages may be in nonconsecutive
2 physical memory space. A completion status structure contains information indicative of
3 the status of transmitted frames. This is implemented as a separate LCP channel. The
4 frame descriptor, which is the first descriptor for every frame, has an optional status field
5 which can be updated after the frame has been transferred to the adapter 80.

6 Referring now to Figure 13, in an example of a transmit LCP channel flow, descriptors
7 1100 are located in the memory 60 of the host 10. Access to the descriptors 1110 and
8 buffers 1110 storing packets 1120 requires address translation through a translation table
9 1130 located in the adapter 80. The buffers 1110 use contiguous space in the virtual
10 address space of the software in the host 10. Each frame 1120 is described by two types
11 of descriptors: a frame descriptor 1140 giving information relating the packet; and, a
12 pointer descriptor 1150 pointing to the buffer 1110 holding the data 1120. Each packet
13 comprises a data payload 1170 preceded by a header 1160 in the same buffer 1180.

14 A write transaction 1190 to the Doorbell updates the number of words 1200 available for
15 use by the adapter 80. This information is stored in the LCP context 140. The transmit
16 LCP context 140 includes a pointer 1210 to the head of the buffer 1110 holding the data
17 to be transmitted. When the LCP channel wins the internal channel arbitration of the
18 ISOC 120, the ISOC 120 reads the descriptors of the LCP channel according to the
19 pointer 1210 in the LCP context 140. Virtual addresses, for both descriptors 1100 and
20 buffers 1110 of the LCP channel, are translated into physical addresses using the
21 translation table 1130 located in the memory subsystem 210 of the adapter 80. The
22 translation table 1130 is updated by the LCP manager 140 during registration of the
23 memory buffers. The ISOC 120 reads the data and frame headers from the buffers 1110
24 into the adapter 80. The frame headers 1160 are then replaced on the ISOC 1320 by a
25 header for the network 30. The packet header and the corresponding data are then
26 transmitted to the network 30.

1 The RX LCP port is used to transfer incoming data from the ISOC 120 to the memory 60
2 used by a software application running on the host 10. TX LCP channels are completely
3 controlled through descriptors initiated by the software on the host 10. RX LCP channels
4 use descriptors from both the software on the host 10 and the ISOC 120. The descriptors
5 initiated by the ISOC 120 are used to control the LCP channel operation to define the
6 destination of a received frame in the memory 60 of the host 10. The descriptors initiated
7 by the software in the host 10 can be used to define the location of buffers where the
8 buffers were not defined through mapping in a translation table. To implement a
9 handshake between the software in the host 10 and the adapter 80, two LCP channels are
10 preferably used: an RX LCP channel for handling the received incoming data structure;
11 and, an RX LCP channel for handling the completion status queue. The completion
12 status is used by the adapter 80 to signal to the software in the host 10 that a frame
13 transfer into the memory 60 of the host 10 is completed. Entries are inserted into the
14 completion queue structure in sequential addresses. Each completion status entry
15 contains a field that is marked by the adapter 80 and pooled by the software in the host 10
16 to check that the entry ownership has been transferred from the adapter 80 to the software
17 in the host 10. One or more RX LCP channels can use the same completion status queue.
18 The sharing of the completion status queue by multiple RX LCP channels is performed by
19 the ISOC 120.

20 An RX LCP channel requires information to indicate the destination address for an
21 incoming packet. The ISOC 120 has two addressing for finding the location of free
22 buffers:

23 Direct addressing mode refers to LCP channels that do not use pointer descriptors
24 to point out a buffer. The destination address is defined either by microcode in
25 the ISOC 120 or read from the context 140.

1 Indirect addressing mode refers to LCP channels that maintain pointers to data
2 buffers in descriptor structures. The descriptors are preferably located in the
3 memory 60 of the host 10.

4 Direct addressing substantially cuts down the latency of processing an incoming packet
5 through the adapter 80. However, it requires registration of memory buffer by the LCP
6 manager 130, including storage of virtual to physical translation information on the
7 adapter 80. The software in the host 10 writes to the channels Doorbell to indicate the
8 amount of words added to the free buffer that can be used by the channel. In direct mode,
9 the following steps are used to determine the address of the destination buffer:

- 10 A. Address A is driven as a command to the LCP engine.
- 11 B. (Optional) Address A is mapped to address A'.
- 12 C. Address A' (if step B is executed) or A (if step B is not executed) is the
13 base address for the destination buffer.

14 In indirect mode, the adapter 80 uses descriptors to find the address of the data buffers.
15 The descriptors are managed by the software in the host 10. The descriptors are
16 preferably located in the memory 60 of the host 10. The term indirect is used to
17 emphasize that the adapter 80 reads additional information to define the destination
18 address. The adapter 80 accesses this information during run-time. Indirect addressing
19 cuts down the amount of the memory n the adapter 80 required to store translation tables.
20 The descriptors are typically located in the memory 60 of the host 10. In indirect mode,
21 the following steps are used to determine the address of the destination buffer:

- 22 A. Address A is driven as a command to the LCP engine.
- 23 B. (Optional) Address A is mapped to address A'.
- 24 C. Address A' (if step B is executed) or A (if step B is not executed) is the
25 address of the pointer descriptor.

- 1 D. The pointer to the buffer, address B, is read from the descriptor.
- 2 E. (Optional) Address B is mapped to address B'.
- 3 F. Address B' (if step E is executed) or B (if step E is not executed) is the
- 4 base address for the destination buffer.

5 Each RX LCP channel uses the following data structures:

6 Access to the Doorbell, implemented as a memory mapped address, informs the
7 adapter 80 of additional data or descriptors available for the adapter 80 to write
8 packet data.

9 An LCP context entry in the memory space 210 of the adapter 80 contains LCP
10 attributes, state, configuration, and status fields.

11 Descriptors pointing to memory buffers for use in indirect mode.

12 A buffer in contiguous virtual address space in the memory 60 of the host 10.

13 A translation table and protection blocks in the memory space 210 of the adapter
14 80 for address mapping.

15 The flow of receiving a packet depends on the following characteristics:

16 Direct or indirect addressing mode.

17 For indirect mode, descriptors are located in the memory 60 of the host 10.

18 For direct mode, address mapping may or may not be used during access to
19 descriptors.

20 Address mapping may or may not be used during access to buffers.

21 For indirect mode, address protection may or may not be used during access to
22 descriptors.

23 Address protection may or may not be used during access to buffers.

1 These characteristics are set for each LCP channel as part of the channel's context 140
2 during the LCP channel initialization.

3 Referring to Figure 14, a flow of receive packets comprises, at step 1300, preparation by
4 software 1310 in the host 10 of free buffer 1320 for the received data. At step 1330, in
5 indirect mode, the software 1310 in the host 10 updates the descriptors 1340. The
6 descriptors 1340 are located in the memory 60 of the host 10. At step 1350, the software
7 in the host 10 rings the Doorbell to notify the adapter 80 of the free buffer space. For
8 indirect mode, the Doorbell provides information indicative of the new descriptors 1340.
9 For direct mode, the Doorbell provides information indicative of added free buffer space.
10 At this stage, the adapter 80 is ready to transfer receive data from the network 30 to the
11 memory 60 of the host 10. Steps 1300, 1330, and 1350 are repeated whenever the
12 software 1310 in the host 10 adds free buffers 1320 to the RX LCP channel. The ISOC
13 120 repeats the following steps for each received packet. At step 1360, the adapter 80
14 receive the data. At step 1370, in indirect mode, the adapter 80 reads descriptors 1340
15 pointing to the location of the free data buffers 1320. At step 1380, data and headers are
16 written into the data buffers 1340. At step 1390, status is updated in the completion
17 queue.

18 Referring to Figure 15, in an example of a receive LCP channel flow, pointer descriptors
19 are not used. Furthermore, no translation tables are used. Data buffers 1400 use
20 contiguous space in the physical address space of software in the host 10 using the buffers
21 1400. Both header and data payload are written to the buffers 1400. A write transaction
22 1410 to the Doorbell updates the data space available for use by the adapter 80. The
23 information is stored in the LCP context 140. The receive/completion LCP context 140
24 includes a pointer 1420 to the head of the buffer 1400 and an offset 1430 to the
25 next/current address used to write new data/completion entries. When the adapter 980
26 receives a packet, it increments the offset 1430 to the next packet location and updates the
27 available data space. A completion entry 1440 is added to a completion LCP 1450 upon

1 completion of a frame reception, upon frame time-out, or for any other frame event that
2 requires awareness from the LCP client 100. The completion entry 1440 contains all the
3 information needed by the LCP client 100 to locate the frame within the LCP data buffer
4 1400. The software in the host 10 uses a field within the completion entry 1440 to
5 recognize that it has been granted ownership of the completion entry 1440.

6 The ISOC 120 allows LCP channels to be used for moving data between the memory
7 subsystem 210 of the adapter 80 and the memory 60 of the host 10. To transfer data from
8 the memory 60 of the host 10 to the adapter 80 a transmit channel is used. To transfer
9 data from the adapter 80 to the memory 60 of the host 10 a receive channel is used.
10 When data is to be transferred from the memory 60 of the host 10 to the adapter 80 a
11 frame descriptor includes a destination address on the bus 340 of the ISOC 120. This
12 address defines the destination of the frame data payload. The packet header is
13 transferred in the usual manner. This allows loading of tables and code into the memory
14 space of the ISOC 120. To transfer data from the memory space of the ISOC 120 to the
15 memory 60 of the host 10 using a receive channel a descriptor is initiated by the RX
16 processor 160. The descriptor include information indicative of both destination address
17 in the memory 60 of the host 10 and source address.

18 Referring now to Figure 16, as indicated earlier, the ISOC 120 comprises RX logic 1500
19 and TX logic 1510. The RX logic 1500 comprises a plurality of registers 1520-1540 for
20 handling interrupts. Likewise, the TX logic 1510 comprises a plurality of registers
21 1550-1570 for handling interrupts. The registers 1550-1570 in the TX logic 1510
22 comprise status registers 1570, processor interrupt mask registers 1560, and host interrupt
23 registers 1550. The registers 1520-1540 in the RX logic 1500 also comprise status
24 registers 1540, processor interrupt mask registers 1530, and host interrupt registers 1520.
25 The registers in the TX logic 1510 and the RX logic 1500 are connected in similar
26 configurations. ISOC level interrupts 1580 are interrupts directed to the host computer
27 system 10. The interrupt line 1580 comprises a logical OR on those bits in the status

1 registers 1570, 1540 which are not masked by the corresponding mask registers
2 1560,1530. These interrupts originate at the following sources: an LCP operation
3 completion 1590; a call from the TX processor 150; a call from the RX processor 160;
4 events detected by the TX logic 1510; and, events detected by the RX logic. Calls from
5 the TX processor 150 and the RX processor 160 are generated by writing to TX and RX
6 call registers respectively. In the TX logic 1510, the mask registers 1560 control passage
7 of interrupts from the corresponding status registers 1540 to the TX processor 150 and the
8 mask registers 1550 control passage of interrupts to the host 10. Similarly, In the RX
9 logic 1500, the mask registers 1530 control passage of interrupts from the corresponding
10 status registers 1540 to the RX processor 160 and the mask registers 1520 control passage
11 of interrupts to the host 10. This arrangement permits the host 10 to acknowledge every
12 possible interrupt generated by an event on the ISOC 120. The arrangement can be
13 employed to interrupt the host 10 for errors where microcode in the ISOC 120 is not
14 expected to be capable of handling event because, for example, the task is too
15 complicated to be handled by the microcode or because the microcode has crashed at the
16 time that error occurs. The arrangement can also employed to call the host 10 as a guard
17 for the microcode. In this case, the microcode is responsible for taking action following
18 the reported error or event. The microcode returns a completion indication to the host 10
19 following handling of the exception.

20 An interrupt to the TX processor 150 is generated as a logical OR on non-masked bits in
21 a status register 1570. This status register is a first level interrupt register that does not
22 define full details of the cause of the interrupt. The TX processor 150 reads the cause of
23 the interrupt from a second level interrupt register. The interrupt is cleared by writing to
24 a clear address in the second level interrupt register. Interrupts to the TX processor 150
25 originate from the following sources: all from the software on the host 10; a call from the
26 RX processor 160; and, events detected by the TX logic 1510. Similarly, an interrupt to
27 the RX processor 160 is generated as a logical OR on non-masked bits in a a status
28 register 1540. This status register is a first level interrupt register that does not define full

1 details of the cause of the interrupt. The RX processor 160 reads the cause of the
2 interrupt from a second level interrupt register. The interrupt is cleared by writing to a
3 clear address in the second level interrupt register. Interrupts to the RX processor 160
4 originate from the following sources: all from the software on the host 10; a call from the
5 TX processor 150; and, events detected by the RX logic 1500.

6 In some embodiments of the present invention, the TX processor 150 may be additionally
7 responsible for handling errors reported by logic shared between the transmit and receive
8 paths. In other embodiments, the RX processor 160 may be responsible for handling such
9 errors from the shared logic

10 In preferred embodiments of the present invention, LCP interrupts 1590 are preprocessed
11 in the interests of reducing processing burden on the software in the host 10. LCP
12 interrupt information is written into the memory 60 of the host 10 to reduce software
13 latency from repeated accesses by the ISOC 120. The generation of a new interrupt
14 indication by each LCP channel is deferred until handling of previous channel interrupts
15 is completed. The processing of LCP interrupts will now be described in detail with
16 reference to Figure 17.

17 Referring now to Figure 17, an interrupt flow between the ISOC 120 and the host 10
18 comprises, at step 1600, either a software application 1610 in the transmission direction
19 or the RX processor 160 in the reception direction setting a CompletionEventRequest bit
20 in a descriptor 1620 for which an interrupt is required. The descriptor 1620 is stored in a
21 descriptor queue 1630. At step 1640, once processing of the descriptor is completed, a
22 completion event indication 1650 is sent to an interrupt FIFO buffer 1660 in the ISOC
23 120 by an interrupt controller of the ISOC 120. An EventMask bit is set in the LCP
24 context 140. Completion event indications are queued in the interrupt FIFO 1660. At
25 step 1670, when preset conditions are met, an Interrupt Control Block (ICB) 1680 is
26 generated by the ISOC 120 from the information stored in the interrupt FIFO 1660. The

1 preset conditions will be described shortly. At step 1690, the ICB 1680 is transferred to
2 the memory 60 of the host 10. ICBs 1680 from the ISOC 120 are stored in a wrapped
3 queue 1700 in the memory 60 of the host 10. At step 1710, an interrupt handler 1720 in
4 the software of the host 10 reads the ICB 1680. At step 1730, the interrupt handler 1720
5 sends the completion event 1650 from the ICB 1680 to the application 1610. At step
6 1750, the application 1610 writes a ClearEventMask bit to the Doorbell register of the
7 LCP channel to enable interrupts from the channel.

8 An active LCP channel can generate one or more completion events 1650 during
9 operation. A completion event 1650 is generated when processing of the descriptor 1620
10 on which the CompletionEventRequest bit is set is completed. The operation of the ISOC
11 120 following a completion event varies depending on the value of the EventMask and
12 the CompletionEvent bit in the context 140. If the EventMask bit is cleared, an indication
13 is sent to the interrupt controller of the ISOC 120 and the EventMask bit is set by the
14 ISOC 120. If the EventMask bit is set and the CompletionEvent bit in the channel's
15 context 140 is cleared, no indication is transferred to the interrupt controller and the
16 CompletionEvent bit is set by the ISOC 120. If the EventMask bit and the
17 CompletionEvent bit in the channel's context 140 are both set, no action is taken. The
18 EventMask bit is cleared at channel initialization. It is also cleared after the
19 ClearEventMask bit is written to the context 140 of the channel via the Doorbell register.
20 If the CompletionEvent in the channel's context 140 is set and the mask bit is cleared by
21 the ClearEventMask bit in the Doorbell register, an indication of the event completion is
22 sent to the interrupt controller and the CompletionEvent bit is cleared. The completion
23 events are logged in the FIFO 1660 by the interrupt controller. Each entry in the FIFO
24 1660 holds a field for describing the number of the LCP channel generating the event.

25 The ICB 1680 is a data structure transferred by the ISOC 120 to the memory 60 of the
26 host 10 via a dedicated LCP channel. Referring to Figure 18, the ICB comprises a header
27 portion and a payload portion. The header portion comprises, at word 0, a status word

1 including an ICB index identifying the ICB 1680, an LCP interrupts valid count
2 indicating the number interrupts in the payload portion, and a time of day (TOD) stamp.
3 The remainder of the ICB 1680 is devoted to the payload portion. The payload portion
4 comprises a plurality of fields each containing the identity of the LCP channel that
5 indicated the completion event. In the example shown in Figure 18, each field in 2 bytes
6 long, and there are 28 fields in the ICB 1680. However, it will be appreciated that, in
7 other embodiments of the present invention, the field size, or the ICB size, or both, may
8 be different.

9 The ICB 1680 is transferred to the memory 60- of the host 10 via a DMA action. The
10 ICB DMA is initiated by any one of the following events:

11 There is at least a predetermined minimum number of event completion
12 indications in the FIFO 1660 and a predetermined minimum time period has
13 passed;

14
15 There is at least one event in the ICB 1680 and a predetermined maximum time
16 period has passed; and,

17
18 The interrupt FIFO buffer 1660 is full.

19 The ICB 1680 is copied to the location in the memory 60 of the host currently pointed to
20 by the context 140 of the interrupt receiver LCP channel. When the ICB write operation
21 is complete, the ISOC 120 asserts an LCP completion bit in its interrupt register. The
22 assertion of the LCP Completion bit generates a maskable interrupt. The LCP
23 Completion bit is cleared by a host read from the interrupt register of the ISOC 120. The
24 ICB LCP channel that moves the ICB 1680 from the ISOC 120 to the memory 60 of the
25 host 10 behaves similarly to other LCP receive channels. Specifically, the contexts 140
26 and buffers associated with the ICB LCP channel are initiated by the LCP manager 130;

1 the buffers used by the ICB LCP channel can be in the same format as other receive LCP
2 channels; and, synchronization between the software on the host 10 and the ISOC is
3 performed by setting new free space words or buffer descriptors through the Doorbell
4 register associated with the channel. In some embodiments of the present invention, the
5 ICB channel differs from other LCP channels in that: the ICB channel does not use the
6 ICB interrupt scheme; completion of an operation on this channel (processing a descriptor
7 or moving a new ICB 1680 to the memory 60 of the host 10) can generate an interrupt;
8 and, the channel is managed through logic in the ISOC 120 rather than a processor in the
9 ISOC 120.

10 The ISOC interrupt handler 1720 on the host 10 reads the interrupt register of the ISOC
11 120. Reading the interrupt register causes a completion of the ICB write operation in the
12 memory 60 of the host 10. LCP completion may be monitored by polling the memory 60
13 for the next index of the ICB 1680 by way of alternative to using the interrupt channel.
14 This is because the ICB has a fixed location and therefore the location of the next ICB is
15 known. The interrupt handler 1740 calls the applications 1610 that handle each one of
16 the channels indicating a completion event 1650. To avoid overhead in ICB processing in
17 the host 10, an LCP channel need not send a complete indication (through the ICB 1680)
18 until the EventMask bit in the channel context 140 is cleared. The EventMask bit is
19 cleared by setting the ClearEventMask bit in the Doorbell write.

20 In some embodiments of the present invention hereinbefore described, the adapter 80 is
21 connected to the CPU 50 and memory 60 of the host computer system 10 via the bus
22 architecture 70. However, in other embodiments of the present invention, the adapter 80
23 may be integrated into the host computer system 10 independently of the bus architecture
24 70. For example, in other embodiment of the present invention, the adapter 80 may be
25 integrated into the host computer system via a memory controller connected to the host
26 memory 60.

1 Additionally, in some embodiments of the present invention hereinbefore described, the
2 adapter 80 was implemented in the form of a pluggable adapter card for insertion into the
3 host computer system 10. It will however be appreciated that different implementation of
4 the adapter 80 are possible in other embodiments of the present invention. For example,
5 the adapter 80 may be located on a mother board of the host computer system, along with
6 the CPU 50 and the memory 60.

7 Variations described for the present invention can be realized in any combination
8 desirable for each particular application. Thus particular limitations, and/or embodiment
9 enhancements described herein, which may have particular advantages to a particular
10 application need not be used for all applications. Also, not all limitations need be
11 implemented in methods, systems and/or apparatus including one or more concepts of the
12 present invention.

13 The present invention can be realized in hardware, software, or a combination of
14 hardware and software. A visualization tool according to the present invention can be
15 realized in a centralized fashion in one computer system, or in a distributed fashion where
16 different elements are spread across several interconnected computer systems. Any kind
17 of computer system - or other apparatus adapted for carrying out the methods and/or
18 functions described herein - is suitable. A typical combination of hardware and software
19 could be a general purpose computer system with a computer program that, when being
20 loaded and executed, controls the computer system such that it carries out the methods
21 described herein. The present invention can also be embedded in a computer program
22 product, which comprises all the features enabling the implementation of the methods
23 described herein, and which - when loaded in a computer system - is able to carry out
24 these methods.

25 Computer program means or computer program in the present context include any
26 expression, in any language, code or notation, of a set of instructions intended to cause a

1 system having an information processing capability to perform a particular function either
2 directly or after conversion to another language, code or notation, and/or reproduction in
3 a different material form.

4 Thus the invention includes an article of manufacture which comprises a computer usable
5 medium having computer readable program code means embodied therein for causing a
6 function described above. The computer readable program code means in the article of
7 manufacture comprises computer readable program code means for causing a computer to
8 effect the steps of a method of this invention. Similarly, the present invention may be
9 implemented as a computer program product comprising a computer usable medium
10 having computer readable program code means embodied therein for causing a a function
11 described above. The computer readable program code means in the computer program
12 product comprising computer readable program code means for causing a computer to
13 effect one or more functions of this invention. Furthermore, the present invention may be
14 implemented as a program storage device readable by machine, tangibly embodying a
15 program of instructions executable by the machine to perform method steps for causing
16 one or more functions of this invention.

17 It is noted that the foregoing has outlined some of the more pertinent objects and
18 embodiments of the present invention. This invention may be used for many
19 applications. Thus, although the description is made for particular arrangements and
20 methods, the intent and concept of the invention is suitable and applicable to other
21 arrangements and applications. It will be clear to those skilled in the art that
22 modifications to the disclosed embodiments can be effected without departing from the
23 spirit and scope of the invention. The described embodiments ought to be construed to be
24 merely illustrative of some of the more prominent features and applications of the
25 invention. Other beneficial results can be realized by applying the disclosed invention in
26 a different manner or modifying the invention in ways known to those familiar with the
27 art.